

Experimental Demonstration of Frequency Regulation by Commercial Buildings – Part II: Results and Performance Evaluation

Evangelos Vrettos^{*1}, Emre C. Kara^{†2}, Jason MacDonald^{‡3},
Göran Andersson^{§1}, and Duncan S. Callaway^{¶4}

¹Power Systems Laboratory, ETH Zurich, Switzerland

²National Accelerator Laboratory (SLAC), California, US

³Grid Integration Group, Lawrence Berkeley National Laboratory (LBNL),
California, US

⁴Energy and Resources Group, University of California, Berkeley, US

May 19, 2016

Abstract

This paper is the second part of a two-part series presenting the results from an experimental demonstration of frequency regulation in a commercial building test facility. In Part I, we developed relevant building models and designed a hierarchical controller for reserve scheduling, building climate control and frequency regulation.

In Part II, we introduce the communication architecture and experiment settings, and present extensive experimental results under frequency regulation. More specifically, we compute the day-ahead reserve capacity of the test facility under different assumptions and conditions. Furthermore, we demonstrate the ability of model predictive control to satisfy comfort constraints under frequency regulation, and show that fan speed control can track the fast-moving RegD signal of the Pennsylvania, Jersey, and Maryland Power Market (PJM) very accurately. In addition, we report the observed effects of frequency regulation on building control and provide suggestions for real-world implementation projects. Our results show that hierarchical control is appropriate for frequency regulation from commercial buildings.

^{*}vrettos@eeh.ee.ethz.ch

[†]ekara@slac.stanford.edu

[‡]jsmacdonald@lbl.gov

[§]andersson@eeh.ee.ethz.ch

[¶]dcal@berkeley.edu

Acronyms

AHU	Air Handling Unit
AS	Ancillary Service
CWS	Central Working Station
HVAC	Heating, Ventilation and Air-Conditioning
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
MPC	Model Predictive Control
PI	Proportional-Integral
PJM	Pennsylvania, Jersey, and Maryland Power Market
RMSE	Root Mean Squared Error
SAT	Supply Air Temperature

1 Introduction

In Part I of this two-part paper, we performed a detailed literature review on theoretical, simulation-based, and experimental work on frequency regulation with commercial buildings. Furthermore, we presented the test facility for our experiment (FLEXLAB), developed relevant building models, and designed a hierarchical controller for reserve scheduling (level 1), building climate Model Predictive Control (MPC) (level 2) and frequency regulation (level 3).

In Section 2 of Part II, we summarize the control and communication architecture, as well as the experiment settings. In Sections 3, 4, and 5 we report extensive experimental results in FLEXLAB for each level of control over a period of one week. We summarize some important findings and suggestions for future work in Section 6, whereas Section 7 concludes.

2 Preparation of the Experiment

2.1 Communication Architecture

We implement the reserve scheduling problem of level 1 (solved once a day) and the MPC of level 2 (solved every 15 minutes) in Matlab, whereas we calculate the fan speed setpoints of level 3 in Python (every 4 seconds) and communicate them




Figure 1: The developed control and communication architecture for building climate control and frequency regulation in the FLEXLAB test facility.

to the Central Working Station (CWS) of FLEXLAB. We used a file-based communication between Python and Matlab based on comma-separated-values (csv) files.

The reserve scheduler stores the computed reserve capacity in the reserve.csv file. A Python script periodically queries the CWS and stores the building measurements in the measure.csv file. Another Python script periodically queries the publicly available database of forecast.io and stores the weather forecasts in the forecast.csv file.¹ The MPC’s feedback from the building is obtained from measure.csv and the weather forecasts from forecast.csv. The optimal air flow rate setpoint calculated by Matlab is stored in the setpoint.csv file. The fan speed setpoint is determined in Python by accessing the setpoint.csv and reserve.csv files, and based on the frequency regulation signal.

Most of the experiment was performed using archived data of the RegD signal from the Pennsylvania, Jersey, and Maryland Power Market (PJM) from December 2012 to January 2013. Although the signal was available with a resolution of 2 seconds, we down-sampled it to 4 seconds due to the expected communication delays. In addition, a connection with PJM was established based on the DNP3 protocol and using a Siemens Jetstream gateway that provided us with the RegD signal in real-time. At the FLEXLAB side, the received data were translated, saved in an SQL database, and pushed by a “RegD signal server” to a “RegD signal client”. The complete communication architecture from PJM to FLEXLAB is graphically shown in Fig. 1. However, network issues at FLEXLAB made the connection unreliable, and therefore we chose to run the live connection with PJM only for one continuous hour.

2.2 Experiment Settings

Since FLEXLAB is not occupied, we emulated the internal heat gains from occupants and equipment using electric heaters as plug loads. The total internal heat

¹Only ambient temperature forecasts are obtained from forecast.io. The solar radiation forecasts are obtained from a clear-sky radiation model, which turned out to be sufficient for the weather conditions during the experiment.




Figure 2: The heater schedule and the actual power consumption. The heat gain is high during working hours and low during non-working hours.




Figure 3: The SAT setpoint and actual values during one week.

gain in both cells was kept lower than the chiller cooling capacity. The heaters' consumption profile was fixed according to the red curve of Fig. 2 using digital timer sockets. The actual heater power (blue curve) fluctuates around the profile due to voltage variations.

Before the start of the experiment, we fixed the manually controlled inlet dampers in the rooms to fully open positions. In addition, we fixed the return air damper to a 100% opening and the outside air damper to a 0% opening, i.e., the return air was fully recirculated. The speeds of primary and secondary chilled water pumps were fixed to 75% and 100% of their rated speeds, respectively. Moreover, we deactivated the existing floor heating system and the heating coil at the Air Handling Unit (AHU).

We set the temperature comfort zone to 21 – 25°C during working hours. An existing Proportional-Integral (PI) controller regulates the Supply Air Temperature (SAT) to 17°C by controlling the position of a cooling valve. The gains of this controller had been tuned for a conventional building operation; therefore, we modified them to achieve a tighter control and reduce the fluctuations of SAT around its setpoint during frequency regulation. The resulting SAT profile during the experiment is shown in Fig. 3. The mean deviation from the SAT setpoint is 0.05°C (there is a small bias to larger SAT values) and the mean absolute deviation is 0.37°C.

Recall that the facility has two building cells with identical construction: cell




Figure 4: The hourly reserve capacities as a percentage of the nominal fan power for 20 November 2015 (left) and 21 November 2015 (right).

1A is used for the frequency regulation experiment, whereas cell 1B serves as a benchmark. Applying the same air flow rate in both cells and recording the temperature, we verified that the two cells are thermally very similar. However, we observed that the same fan speed setpoint induces a slightly different air flow rate in the two cells due to small differences in the AHUs. To compensate for this, we fitted different fan models for the two cells (the parameters for cell 1A are given in [1, Table IV]).

The electricity cost was assumed equal to $c_k = 0.18$ €/kWh, whereas the reserve capacity payment was fixed to a 10% higher value, i.e., $\lambda_k = 0.198$ €/kWh. The goal of this experiment is to demonstrate the technical feasibility of reserve provision from commercial buildings; therefore, we chose a relatively high capacity payment to incentivize reserve provision.

2.3 Experiment Plan

The experiment was organized into two parts. The first part took place from 15 to 18 November 2015 and relied on an “older” building model identified with data from June-July 2015 (see [1, Table II]). On 19 November the experiment was paused and a new building model was identified using the recently collected data (see [1, Table III]), which was used in the second part of the experiment from 20 to 21 November.

3 Reserve Scheduling (Level 1)

In this section, we present results relevant to the reserve scheduler. Two main factors that determine the amount of reserves are the building’s energy capacity and the symmetry of reserve capacity. Apart from the physical properties of the building, the energy capacity depends also on the comfort zone’s width. In this experiment, we specifically address the effect of enlarging the comfort zone during unoccupied hours to $12 - 35^\circ\text{C}$ (the so-called night setback). We performed six full-day experiments with symmetric (equal up- and down-reserves) or asymmetric




Figure 5: The hourly reserve capacities as a percentage of the nominal fan power for 15 November 2015 (left) and 17 November 2015 (right).

reserves, and with or without night setback. Note that the same price is assumed for up- and down-reserves in the asymmetric case.

Fig. 4 shows the results for 20 November when symmetric reserve capacities were assumed and for 21 November when asymmetric capacities were used (in both days night setback was applied). The capacities are reported in % of the fan rated power (2500 W). The reserve capacity is maximized at night when the comfort zone is enlarged, and during the hottest part of daytime. In case of symmetric reserves, the maximum capacity is slightly less than 40% of the rated fan power. In case of asymmetric reserves, the maximum up-reserve capacity is approximately 15%, whereas the down-reserve capacity is more than 60%. These experimental results are in agreement with relevant simulation results in [2,3], and show that down-reserves (consumption increase) are preferable for commercial buildings equipped with energy-efficient controllers, because down-reserves can be provided without increasing the baseline consumption and the energy cost.

Fig. 5 compares the experimental results of 15 November when setback was used with those of 17 November when no setback was applied (the reserve capacities were symmetric in both dates). With setback most reserve is provided at night, whereas without setback the reserve provision coincides with the highest cooling load in the middle of the day. Although the experiment was conducted with setback and symmetric reserves both on 15 and 20 November, the capacity profiles during daytime are considerably different due to different weather conditions and building models.

To have a fair comparison under the same external conditions, we simulated the reserve capacity scheduling for all combinations of symmetry and setback using the building model. The simulation and experimental results are shown in Table 1. The capacity ranges from low values below 1% to high values nearly 50%, and it heavily depends on reserve symmetry, setback, and weather conditions. The night setback increases the capacity by 177.0% on average for symmetric reserves, by 107.0% for asymmetric up-reserves, and by 150.7% for asymmetric down-reserves. If setback is already used, adopting asymmetric capacities instead of symmetric capacities reduces the up-reserves by 55.6% but increases the

Table 1: Experimental (bold) and simulated (normal font) daily average reserve capacities in % of fan nominal power (2500 W).

Date	Symmetric, Setback		Asymmetric, Setback		Symmetric, No setback		Asymmetric, No setback	
	R_u	R_d	R_u	R_d	R_u	R_d	R_u	R_d
15/11	15.61	15.61	6.45	26.15	1.91	1.91	1.64	4.42
16/11	9.09	9.09	3.90	14.88	0.74	0.74	0.90	1.85
17/11	11.44	11.44	5.39	21.21	2.24	2.24	1.85	4.59
18/11	16.70	16.70	7.78	31.75	3.94	3.94	3.28	9.91
20/11	28.95	28.95	11.89	49.66	15.07	15.07	7.82	28.81
21/11	22.10	22.10	10.72	46.55	13.60	13.60	6.79	26.30

down-reserves by 83.1%, and so the net effect is an increase of 13.7% in the total capacity.

4 Room Climate Control (Level 2)

4.1 Comfort Satisfaction

Experimental results for 17 November are shown in Fig. 6, where the top plot shows the temperature trajectories in cells 1A and 1B, the middle plot presents the forecasts and actual values for ambient temperature and solar irradiance², whereas the bottom plot shows the SAT and the air flow rate in cell 1A. The comfort zone is indicated with red: the actual upper limit (red solid line) is 25°C, but a tighter limit of 24°C (red dashed line) is used within the MPC to account for modeling and prediction errors. Similar results for 18, 20, and 21 November are shown in Figs. 7, 8, and 10, but without including the SAT and air flow rate plot due to space limitations.

In Figs. 6 and 7 the cell 1B is under energy efficient operation and the temperature remains close to the upper limit of the comfort zone. On the other hand, in Figs. 8 and 10 the cell 1B is in a “regulation-ready” operation mode, namely the consumption of the Heating, Ventilation and Air-Conditioning (HVAC) system is scheduled identically to cell 1A to allow reserve provision, but no regulation signal is received. For this reason, the temperature trajectories of the two cells are very close to each other for most of 20 and 21 November.³

²The total global irradiance is shown, which includes the long-wave radiation losses from the building envelope to the atmosphere, and it can be negative at night. This effect is known as nighttime radiation cooling [4].

³The discrepancies from 12.00 to 19.00 on 20 November are due to the calibration differences between the fan models of the two cells (see Section 2.2). The discrepancies from 07.00 to 17.00 on 21 November are because of interruptions in the hierarchical control in cell 1B due to server connection timeout error from approximately 07.00 to 11.00. When the server was unresponsive,

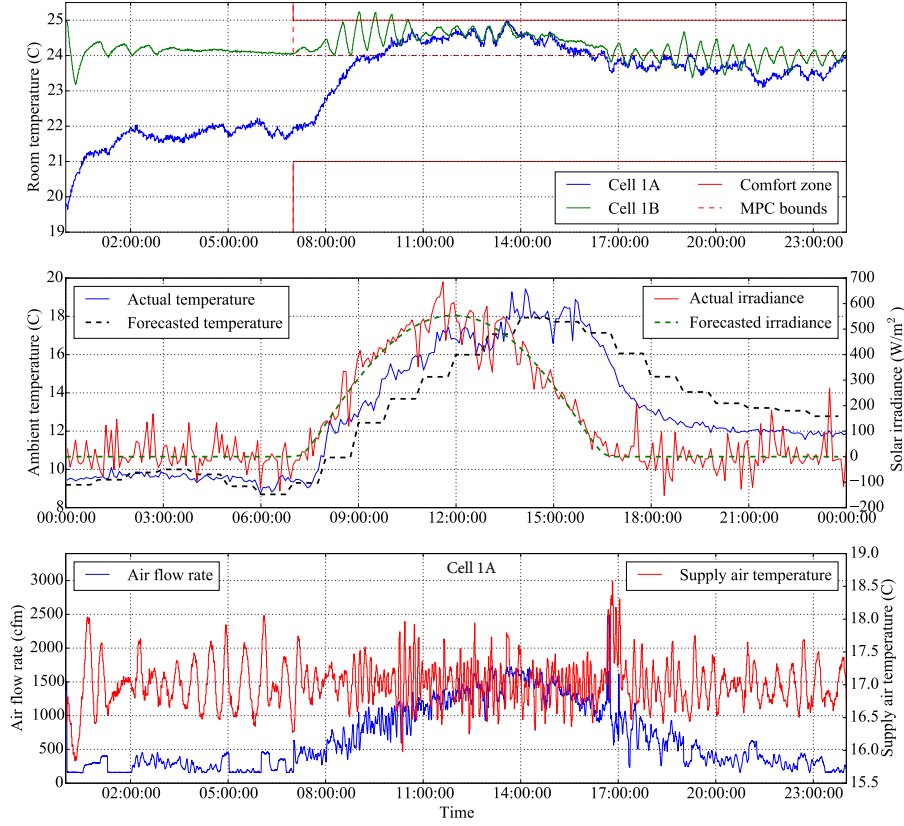


Figure 6: Experimental results for the room climate controller under frequency regulation on 17 November (symmetric reserves, without night setback).

The temperature trajectory of cell 1B in Figs. 6 and 7 remains mostly in the band $[24 - 25]^{\circ}\text{C}$, which illustrates the necessity of tightening the comfort zone constraints in the MPC to compensate for modeling errors. The temperature trajectory of cell 1A is more variable and it follows the scheduled reserve and air flow rate. On 17 November (Fig. 6) frequency regulation is provided while respecting the comfort zone.

However, on 18 November (Fig. 7) the comfort zone is violated from 13.00 to 16.00 in cell 1A, but not in cell 1B. This happens because: (i) the ambient temperature is higher than the day-ahead forecast from the beginning of the day until 15.00, and (ii) asymmetric reserves are used. The asymmetry allows for a more aggressive scheduling with a larger down-reserve capacity on 18 November in comparison with 17 November when symmetric reserves are used (see Table 1).

The control performance is significantly better on 20 and 21 November (Figs. 8 and 10) despite the large discrepancies between the day-ahead ambient temperature forecasts and the actual values. The improvement is due to the recently calibrated the cell was controlled by an existing fallback controller.




Figure 7: Experimental results for the room climate controller under frequency regulation on 18 November (asymmetric reserves, without night setback).

building model (see Section 2.3). No comfort zone violations occur and moreover the temperature is below the MPC constraint of 24°C for most of the time. Therefore, periodic model calibration (for example on a weekly or daily basis) is important to account for seasonality and eliminate systematic errors.

These results show that if the model and weather predictions are sufficiently accurate, the robust reserve scheduler allows a commercial building to bid in day-ahead markets for frequency regulation. On the other hand, if the modeling and prediction errors exceed the controller’s robustness margin, reserve provision for frequency regulation might have an adverse effect on occupant comfort.

The temperature trajectory in Figs. 8 and 10 is typical for a building with a night setback. The controller chooses to overcool the space at night in order to generate higher revenue by offering a larger reserve capacity. In contrast, the reserve capacity is smaller during working hours and the room temperature is higher. A comparison of the temperature trajectories in cells 1A and 1B shows that tracking the RegD signal has little effect on room temperature due to the signal’s limited energy content.

4.2 Model and Estimator Performance

Fig. 9 compares the out-of-sample performance of the older model (left plot) and the new model (right plot). The blue curve is the measured room temperature, the green curve is the estimated temperature with the Kalman filter, the red curve corresponds to a day-ahead model prediction, whereas the orange curve shows the step-ahead temperature predictions. Clearly, the new model outperforms the older




Figure 8: Experimental results for the room climate controller under frequency regulation on 20 November (symmetric reserves, with night setback).




Figure 9: Model and Kalman filter performance. Left: Results for 17 November with the older model. Right: Results for 20 November with the new model.

one, especially for the day-ahead predictions. This is why the performance of the level 2 controller is much better on 20 November than on 17 November in terms of comfort zone violations.

The effect of model accuracy on MPC operation is shown in Fig. 11. On 17 November the model mismatch is large, which results in a significant discrepancy in the scheduled air flow rate and fan power between levels 1 and 2. The MPC reacts on the modeling error by reducing the cooling power in level 2 during night hours and increasing it during daytime. In this way, the MPC provides the same amount of electric reserve in daytime with less change in air flow rate by taking advantage of the nonlinear fan curve. On the other hand, the model mismatch is small on 20 November and so the air flow and fan power schedules of level 1 and level 2 are similar. In fact, level 2 consistently schedules less cooling power than level 1 because the air flow constraints are relaxed [1, Equations 15, 24], and the




Figure 10: Experimental results for the room climate controller under frequency regulation on 21 November (asymmetric reserves, with night setback).




Figure 11: Left: Air flow rate schedule in level 1 and level 2. Right: Fan power schedule in level 1 and level 2.

reserve scheduling in level 1 is robust and thus conservative.

4.3 Fan Heat Gain at High Speeds

We present results on the dependence of SAT and cooling valve opening on fan speed in Fig. 12, where the blue points are measurements and the red trend is a polynomial fit on them. As expected, the trend in cooling valve opening is increasing because the higher the fan speed the more cooling is required from the chilled water loop. The trend in SAT is a flat line for fan speeds up to 50%. However, for speeds above 50% (and especially above 70%) there is a clear increasing SAT trend despite the increased cooling valve opening.

These results lead to an interesting observation: the heat gain due to the rotation of the fan is significant at high speeds and it cannot be effectively rejected by ex-




Figure 12: The dependence of SAT and cooling valve opening on fan speed.

changing heat with the chilled water loop. According to Fig. 12, if the fan operates at a 70% speed or higher, the SAT will likely have a steady-state deviation from the setpoint 17°C that can be as high as 1°C. Steady-state SAT deviations might result in comfort zone violations, because the controller assumes the SAT fixed to 17°C. This did not create problems in our experiment because the scheduled fan speed by the MPC was at most 70%.

4.4 Effect on Energy Consumption

A major concern when providing Ancillary Services (AS) with commercial buildings is the effect on energy consumption. Ref. [5] reported a round-trip efficiency of 46% when a building responded to demand response events in an experiment. There are two types of efficiency losses relevant to frequency regulation: “reserve availability efficiency loss” and “reserve utilization efficiency loss” [6]. The first one is the efficiency loss due to scheduling the consumption in an energy suboptimal way to be able to provide frequency reserves, if requested. The second one is the additional efficiency loss while tracking the frequency regulation signal.

We report efficiency results in Table 2 for: (i) 15-18 November when the cell 1B was under energy efficient operation (to quantify the reserve availability efficiency loss); and (ii) 20 November when cell 1B was in regulation-ready operation mode (to quantify the reserve utilization efficiency loss). The efficiency loss is calculated comparing the energy consumption of cell 1A with that of the benchmark cell 1B. We use two different definitions of energy consumption: (i) electric energy consumption of the fan, and (ii) thermal cooling power consumption of each cell. The latter is calculated based on the chilled water flow rate (\dot{m}_{cw}), as well as the supply ($T_{ch,s}$) and return ($T_{ch,r}$) chilled water temperatures using

$$P_{cool} = \dot{m}_{cw} \cdot (T_{ch,r} - T_{ch,s}) . \quad (1)$$

Based on the results of Table 2, the reserve availability efficiency loss is significant and equal to approximately 68% in terms of fan consumption and 11% in terms of cooling power from the chiller. However, the additional consumption in cell 1A is not entirely wasted because it results in a lower average temperature. When cell 1B is in the regulation-ready mode, the cell 1A consumes less energy

Table 2: Effect of frequency regulation on energy consumption

1B operation mode	Energy efficient		Regulation-ready			
	15/11-18/11		20/11 (0-24 h)		20/11 (0-12 h)	
Cell	1A	1B	1A	1B	1A	1B
Fan energy (kWh)	27.22	16.23	19.24	20.91	9.40	10.19
Cooling (gpm · F)	1989.92	1800.55	772.16	788.69	328.78	331.09
Mean temp. (°C)	22.85	24.43	21.45	21.26	21.05	21.05




Figure 13: The average MPC computation time depending on the time of the day and on reserve symmetry.

than cell 1B despite frequency regulation. The non-negligible difference in the average temperature of the two cells is due to the imperfections in fan model calibration and the limited temperature sensor accuracy. However, even from 00.00 to 12.00 when both average cell temperatures are 21.05°C, the consumption of cell 1A is still lower than that of cell 1B. This result indicates that the “reserve utilization efficiency loss” is negligible while tracking a fast-moving regulation signal like RegD.

4.5 MPC Computation Time

The MPC computation time is sufficiently low for our demonstration. As shown in Fig. 13, the longest computation time is 150 s for symmetric reserves and 65 s for asymmetric reserves. The computation time for asymmetric reserves is lower because the problem is simpler and smaller [1, Section IV.D].

The computation time decreases at the end of the day because a reducing MPC horizon is used. After the 70th time step, when the MPC prediction horizon is smaller than 26 time steps (6.5 hours), the computation time is less than 2 s. Therefore, Fig. 13 can be used to select the prediction horizon’s length depending on the maximum allowable computation time. Since the computation time grows exponentially with the number of variables of the nonlinear optimization problem, a shorter prediction horizon might be necessary for larger buildings.

5 Regulation Signal Tracking (Level 3)

5.1 Control Performance Metrics

In this section, we present results from level 3 and evaluate the tracking performance of the regulation signal. The following metrics are used

$$e_{t,k} = e_{c,k}/P_{d,k}, \quad e_{c,k} = P_{d,k} - P_{f,k} \quad (2)$$

$$e_{r,k} = \begin{cases} e_{c,k}/R_{u,k}, & \text{if } w_k < 0 \\ e_{c,k}/R_{d,k}, & \text{if } w_k \geq 0 \end{cases} \quad (3)$$

$$e_{me} = (1/N_{exp}) \cdot \sum_{k=0}^{N_{exp}-1} e_{c,k} \quad (4)$$

$$e_{mae} = (1/N_{exp}) \cdot \sum_{k=0}^{N_{exp}-1} |e_{c,k}| \quad (5)$$

$$e_{rmse} = \sqrt{(1/N_{exp}) \cdot \sum_{k=0}^{N_{exp}-1} e_{c,k}^2} \quad (6)$$

$$e_{t,mape} = (1/N_{exp}) \cdot \sum_{k=0}^{N_{exp}-1} |e_{t,k}| \quad (7)$$

$$e_{r,mape} = (1/N_{exp}) \cdot \sum_{k=0}^{N_{exp}-1} |e_{r,k}|, \quad (8)$$

where $P_{f,k}$ denotes the instantaneous fan power, w_k denotes the normalized regulation signal, and N_{exp} denotes the experiment duration. The metrics $e_{t,k}$ and $e_{r,k}$ are relative instantaneous errors but the normalization is performed using the desired fan power $P_{d,k}$ in $e_{t,k}$, and the up- ($R_{u,k}$) or down-reserve capacity ($R_{d,k}$) in $e_{r,k}$. The mean error e_{me} is used to measure any biases in the control response, whereas e_{mae} is the Mean Absolute Error (MAE) during the experiment. The Root Mean Squared Error (RMSE) e_{rmse} penalizes more the large control errors, for example due to overshoots and undershoots. The metric $e_{t,mape}$ is the tracking Mean Absolute Percentage Error (MAPE), and $e_{r,mape}$ is the reserve MAPE. We use the metric $e_{r,mape}$ because it describes the relative size of the control error with respect to the reserve capacity.

In addition, we use the score proposed by PJM for evaluating the performance of frequency regulation. The total score S_{tot} consists of three parts, namely the correlation score S_c , the delay score S_d and the precision score S_p , which are defined as [7]

$$S_c = \max_{\tau \in [0, 5 \text{ min}]} (R_{cor}) \quad (9)$$

$$S_d = \left| \frac{\tau^* - 5 \text{ min}}{5 \text{ min}} \right|, \quad \tau^* = \underset{\tau \in [0, 5 \text{ min}]}{\operatorname{argmax}} (R_{cor}) \quad (10)$$

$$S_p = 1 - (1/n) \cdot \sum_{k=1}^n \left| \frac{e_{c,k}}{\bar{P}_{d,h}} \right| \quad (11)$$

$$S_{tot} = (1/3) \cdot S_c + (1/3) \cdot S_d + (1/3) \cdot S_p. \quad (12)$$

The correlation score is the maximum correlation R_{cor} of the desired power $P_{\text{d},k}$ and fan power $P_{\text{f},k}$, and τ^* is the time shift at which the correlation is maximized (τ takes a value from 0 to 5 min with a step of 10 s). We calculate the delay score based on the time shift with maximum correlation. In the precision score calculation, we normalize the absolute control error by the average hourly value of the reference signal $\bar{P}_{\text{d},h}$, whereas the total score is a weighted sum of the individual scores.

5.2 Experimental Time Series Results

In Fig. 14 we present results from the operation of level 3 controller from 18:30 to 19:30 on 20 November 2015. The duct pressure is quadratic to the fan speed, as expected from the fan laws. Since the duct system is designed to sustain the pressure corresponding to the maximum fan speed, and because the fan speed does not exceed its maximum value (90%) during frequency regulation, pressure constraints were not necessary in the reserve scheduling and MPC formulations in our experiment.

The RegD signal changes direction very often and has a limited energy content. During periods of time when the RegD signal is relatively flat, or the reserve capacity is low, the PI controller is active. On the other hand, whenever the changes in fan power are rapid, the control switches to the model-based feedforward control.

The tracking of the RegD signal is generally very good. However, when large rapid changes in fan power are requested, overshoots or undershoots might appear. In addition, if the reserve capacities change significantly at the beginning of each full hour, temporarily large errors might occur. In general, the instantaneous percentage errors $e_{\text{t},k}$ and $e_{\text{r},k}$ are higher at a low operating fan power and low reserve capacity.

5.3 Evaluation of Tracking Performance

The performance metrics (2)-(8) for the 6 days of Table 1 are presented in Table 3. $e_{\text{r},\text{mape}}$ is larger than $e_{\text{t},\text{mape}}$ because small reserve capacities are offered for a large part of the experiment. The mean error e_{me} has a negative bias, which means that the fan power is more often higher than the desired setpoint because the control overshoots are larger than the undershoots.

We investigate the dependence of control performance on the minimum reserve capacity, which we call “reserve threshold” and denote by R_{thr} . The metrics e_{mae} , e_{rmse} , $e_{\text{t},\text{mape}}$ and $e_{\text{r},\text{mape}}$ are recalculated considering only the time steps when $R_{\text{u},k} \geq R_{\text{thr}}$ if $w_k < 0$, and $R_{\text{d},k} \geq R_{\text{thr}}$ if $w_k \geq 0$. We repeat this procedure for different R_{thr} values and present the results in Fig. 15.

In contrast to $e_{\text{t},\text{mape}}$, $e_{\text{r},\text{mape}}$ decreases rapidly as R_{thr} increases in the range $[0, 200]$ W. This happens because (for the same absolute control error) $e_{\text{r},k}$ decreases if $R_{\text{u},k}$ or $R_{\text{d},k}$ increase. On the other hand, e_{mae} and e_{rmse} generally increase as R_{thr} increases because the higher the reserve capacity the larger the fan




Figure 14: Fan control and RegD signal tracking for a period of 1 hour.

power change, and thus the higher the errors due to the overshoots and undershoots. Fig. 15 can provide us with bounds on the reserve capacity from a tracking performance point of view.

In Table 4 we report the PJM scores calculated separately for the period 15 - 18 November and the period 20 - 21 November. Different scores are calculated for each hour (only if the reserve capacity is non-zero) [7], and the values in Table 4 are hourly averages. The frequency regulation performance is exceptional during the whole experiment. For comparison, the highest possible total score is $S_{\text{tot}} = 1$ and the minimum S_{tot} accepted by PJM is 0.75. The scores are slightly higher for 20 - 21 November because the building provides a larger reserve capacity compared with 15 - 18 November.




Figure 15: Dependence of tracking and reserve errors on the reserve threshold.

Table 3: RegD tracking performance metrics during the experiment

Metric	e_{me}	e_{mae}	e_{rmse}	$e_{t,mape}$	$e_{r,mape}$
Value	-5.66 W	12.45 W	27.00 W	3.58%	8.23%

5.4 Effect on Supply Air Temperature

Tracking the fast-moving RegD signal introduces high frequency oscillations on SAT as shown in Fig. 16. The SAT of cell 1A (frequency regulation) oscillates more than that of 1B (regulation-ready mode), especially after sudden changes in the regulation signal that induce sudden changes in the air flow rate. In addition, large excursions in SAT occur in both cells when the MPC changes the air flow setpoint significantly, for example at hour 08.00. Moreover, the magnitude of SAT oscillations is high at low air flow rates, for example from 05.00 to 08.00.

5.5 Effect of Fan Control on Chiller Power

The fan and the chiller are thermally coupled through the chilled water loop, hence, it is worth investigating if the operation of the chiller is affected while providing frequency regulation with the fan. In Fig. 17 we present relevant experimental results for a duration of 10 hours. The top plot shows the instantaneous and hourly-average electric power of the fan in cell 1A and the chiller. The bottom plot shows the cooling power in the chilled water loop for cells 1A and 1B calculated with (1).

The chiller has two stages and the electric power consumption is relatively constant at each stage. The chiller's cycling depends on the cooling load, which in turn depends on the fan power and ambient conditions. In general, as the fan power increases the chiller cycles more often and remains longer at the on state. This is shown in Fig. 17 where the average chiller power (green line) generally follows the average fan power (black line).

The effect of regulation is visible on the cooling power that fluctuates more in cell 1A compared with cell 1B (regulation-ready mode). Whenever the fan power increases, the cooling load also increases and the SAT tends to decrease. This is sensed by the SAT controller that opens the cooling valve to compensate for the SAT decrease, which in turn increases the cooling power in the chilled water loop.

Table 4: PJM scores for tracking the RegD signal

Score	S_c	S_d	S_p	S_{tot}
15 – 18 Nov.	0.89	0.97	0.96	0.94
20 – 21 Nov.	0.96	0.99	0.98	0.98




Figure 16: Air flow rate and SAT in cells 1A and 1B on 20 November.

The delay in cooling power response depends on the time constant of the cooling valve’s controller.

Despite the oscillations in cooling power, there is no observable effect on chiller’s cycling and electric power. This happens because: (i) the chilled water is stored in a tank that provides some inertia; and (ii) the RegD signal is approximately zero-mean. Note that the gradual reduction in the hourly-average chiller electric power from 19.00 to 00.00 in Fig. 17 is mainly the result of a lower cooling need due to ambient temperature drop, rather than a side-effect of frequency regulation.

These results indicate that frequency regulation can be provided with fan control without side-effects on chiller consumption. However, this does not necessarily hold for regulation signals with a larger energy content such as RegA. In addition, chillers with continuous rather than duty-cycle control will likely display a more observable impact on consumption while providing frequency regulation, especially if there is no chilled water storage tank. In these cases, the level 3 controller should be revised, which is an interesting direction for future work.

5.6 Analysis of Communication Delays

A challenge in this experiment was the communication delays in measurements and actuation, which result in the overshoots and undershoots in fan power in Fig. 14. In Fig. 18 we present a histogram of the experienced delays during the whole experiment. The probability distribution of the delay is positively skewed with a mean value of 2.89 s and a 95%-percentile of 2.99 s. In fact, there exist a few very large delays in excess of 5 s due to temporary unresponsiveness of the CWS, which are




Figure 17: Effect of frequency regulation on the chiller and cooling power.




Figure 18: Histogram of communication delays during the experiment.

not included in Fig. 18. Despite the fact that the average delay is large compared with the time step of level 3 controller (4 s), the tracking performance of RegD signal is very good.

6 Lessons Learned and Outlook

6.1 Lessons Learned

Hierarchical control is an efficient way to provide frequency regulation with commercial buildings because time-separated tasks are considered individually. Three control layers are essential: (i) a reserve capacity scheduler, (ii) a building climate controller to satisfy comfort while leaving enough slack for reserves, and (iii) a controller to track the regulation signal.

Frequency regulation accuracy: High-quality frequency regulation can be provided by fan speed control. The RegD signal tracking is excellent even with large communication delays in the building automation system. A switched con-

troller comprised of a feedforward controller and a PI feedback controller with gain scheduling provides a fast response without compromising stability. This results in a total PJM score as high as 0.98, which is well above PJM's limit of 0.75.

Means to increase reserve capacity: In our experiment, the fan provided 0.74 – 49.66% of its rated power as reserve capacity, depending on ambient conditions and reserve assumptions. Allowing asymmetric reserve capacities and using a night setback are effective ways to increase the reserve potential from commercial buildings. In fact, down-reserves are preferable for buildings because the capacity can be offered without increasing baseline energy consumption.

Occupant comfort: If the building bids in day-ahead AS markets, respecting occupant comfort might be challenging if the building model and weather forecasts are not very accurate. Furthermore, asymmetric reserves result in a more aggressive scheduling that might increase comfort zone violations.

Building model: An accurate building thermal model is essential for comfort satisfaction, especially in day-ahead AS markets. Periodic calibration of the building model helps to account for seasonality and eliminate offsets in modeling error.

Advantages of MPC: Perhaps the most important advantage of MPC is that it identifies the optimal balance between reserve provision and energy efficiency. MPC additionally provides us with a baseline consumption ahead of real-time operation, which is beneficial from a practical point of view. Moreover, due to its predictive closed-loop nature it reacts to modeling and weather prediction errors in a way that minimizes occupant discomfort.

Robustness measures: It is important to consider the regulation signal uncertainty when scheduling the reserve capacity. A conservative modeling of this uncertainty builds robustness to weather prediction and building modeling errors. Additional robustness can be obtained by tightening the comfort zone constraints in the MPC, and allowing a larger fan speed control band in the MPC compared with the reserve scheduler.

Effects of frequency regulation on building control: Frequency regulation might introduce oscillations in SAT that can be reduced by appropriately tuning the cooling valve controller. In addition, if the MPC schedules the fan speed at very high values, the cooling loop might not be able to reject the additional heat gain due to fan rotation. On the positive side, there is little effect on the average energy consumption of the chiller while tracking an energy-constrained frequency regulation signal by controlling the fan power. However, the impact of fan control on chiller cycling may prevent the building from accurately following the regulation signal when measured against a baseline that includes the combined consumption of the fan and the chiller. This is an interesting area of further study.

Energy consumption: Provision of frequency reserves entails some efficiency loss. The efficiency loss due to scheduling the HVAC consumption in a suboptimal way compared with an energy efficient building control can be as high as 67%. On the other hand, the efficiency loss while tracking frequency regulation signals with limited energy content is negligible.

6.2 Outlook

There are several avenues for follow-up work. Two direct extensions are to repeat the experiment with the RegA signal of PJM, which is slower but has more energy content, and/or with the heating loop of the AHU enabled. In addition, performing the frequency regulation experiment using all four buildings of FLEXLAB will leverage the full potential of hierarchical control and verify the scalability of the approach.

In some HVAC systems a duct pressure controller regulates the pressure to a fixed setpoint. The combined operation of this controller and the dampers of each zone might reject the frequency regulation action [8]. This is an important challenge that could not be addressed in this experiment at FLEXLAB as it requires testing on a large building.

The reserve scheduling optimization problem might be computationally heavy for buildings with many zones. An alternative is to approximate the nonlinear fan power curve with a piecewise affine function by introducing binary variables. The bilinear building dynamics can be approximated with sequential convex optimization [9], but the convergence is not guaranteed. Finally, the conservatism of the reserve scheduling problem can be reduced by generating scenarios from historical frequency regulation signals at the cost of reducing robustness.

7 Conclusion

In Part II of this two-part paper, we reported experimental results for frequency regulation from a commercial building test facility (FLEXLAB). The results are very encouraging: the test building can track fast-moving signals such as RegD reliably, with very high accuracy, and with minimal effect on occupant comfort and the operation of the HVAC system. The results also indicate that a hierarchical control approach is appropriate for frequency regulation with day-ahead bidding of the reserve capacity, and it can be used in field tests and real-world implementations in larger buildings.

References

- [1] E. Vrettos, E. C. Kara, J. MacDonald, G. Andersson, D. Callaway, Experimental demonstration of frequency regulation by commercial buildings – Part I: Modeling and hierarchical control design, available at arXiv.org.
- [2] E. Vrettos, F. Oldewurtel, F. Zhu, G. Andersson, Robust provision of frequency reserves by office building aggregations, in: World Congress of the International Federation of Automatic Control (IFAC), Cape Town, South Africa, 2014.

- [3] E. Vrettos, F. Oldewurtel, G. Andersson, Robust energy-constrained frequency reserves from aggregations of commercial buildings, *IEEE Transactions on Power Systems* (early access).
- [4] B. Givoni, *Passive low energy cooling of buildings*, John Wiley & Sons, 1994.
- [5] I. Biel, I. Hiskens, S. Backhaus, Round-trip efficiency of fast demand response in a large commercial air conditioner, *Energy and Buildings* 97 (2015) 47–55.
- [6] E. Vrettos, G. Andersson, Scheduling and provision of secondary frequency reserves by aggregations of commercial buildings, *IEEE Transactions on Sustainable Energy* 7 (2) (2016) 850–864.
- [7] PJM manual 12: Balancing operations, Available online at www.pjm.com/~media/documents/manuals/m12.ashx, last accessed 22.02.2016.
- [8] P. Zhao, G. P. Henze, S. Plamp, V. J. Cushing, Evaluation of commercial building HVAC systems as frequency regulation providers, *Energy and Buildings* 67 (2013) 225–235.
- [9] F. Oldewurtel, Stochastic model predictive control for energy efficient building climate control, Ph.D. thesis, ETH Zürich (2011).